arXiv:1512.04188v1 [cs.DS] 14 Dec 2015

Hypergraph Two-coloring in the Streaming Model*

Jaikumar Radhakrishnan† Saswata Shannigrahi ‡ Rakesh Venkat §

Abstract

We consider space-efficient algorithms for two-coloring n-uniform hypergraphs $H = (V, E)$ in the streaming model, when the hyperedges arrive one at a time. It is known that any such hypergraph with at most $0.7\sqrt{\frac{n}{\ln n}}2^n$ hyperedges has a two-coloring [4], which can be found deterministically in polynomial time, if allowed full access to the input.

- Let $s^D(v, q, n)$ be the minimum space used by a deterministic (one-pass) streaming algorithm that on receiving an n-uniform hypergraph H on v vertices and q hyperedges ($q \leq 0.7\sqrt{\frac{n}{\ln n}}2^n$), produces a proper two-coloring of H. We show that $s^D(n^2, q, n) = \Omega(q/n)$.
- Let $s^R(v, q, n)$ be the minimum space used by a randomized (one-pass) streaming algorithm that on receiving an n-uniform hypergraph H on v vertices and q hyperedges with high probability produces a proper two-coloring of H (or declares failure). We show that $s^R(v, \frac{1}{10}\sqrt{\frac{n}{\ln n}}2^n, n) = O(v \log v)$.
- We show that for any $4 \leq t \leq \frac{n^2}{2n-1}$, $s^R\left(\frac{n^2}{t}, 2^{n-2}\exp(\frac{t}{8}), n\right) = O(n^2/t)$; in particular, this shows that every n-uniform hypergraph with at most $\frac{n^2}{t}$ vertices and $2^{n-1}\exp(\frac{t}{8})$ hyperedges is two-colorable.

The above results are inspired by the study of the number $q(n)$, the minimum possible number of hyperedges in a n-uniform hypergraph that is *not* two-colorable. It is known that $q(n) = \Omega(\sqrt{\frac{n}{\ln n}}2^n)$ and $q(n) = O(n^2 2^n)$. The lower bound (due to Radhakrishnan and Srinivasan [4]) has a corresponding algorithm to deterministically produce the two-coloring; the upper bound due to Erdős is obtained by picking about $n^2 2^n$ hyperedges randomly from a vertex set of n^2 vertices. Our first result shows that no efficient deterministic streaming algorithm can match the performance of the algorithm in [4]; the second result shows that there is, however, an efficient randomized algorithm for the task; the third result shows that that if the number of vertices is substantially smaller than n^2, then every non-two-colorable hypergraph has significantly more than $n^2 2^n$ hyperedges.

1 Introduction

A hypergraph $H = (V, E)$ is a set system E defined on a finite universe V (called the vertex set). The sets in the set system are called *hyperedges*. We consider n-uniform hypergraphs, that is, hypergraphs whose hyperedges all have n elements. We say that a hypergraph is two-colorable, or has Property B, if there is an assignment of colors $\chi : V \to \{\text{red}, \text{blue}\}$ to the vertex set V such that

*A part of this work was reported earlier in the conference paper, *Streaming Algorithms for 2-Coloring Uniform Hypergraphs* by Jaikumar Radhakrishnan and Saswata Shannigrahi, WADS (2011): 667-678.

†Tata Institute of Fundamental Research, Mumbai. E-mail: jaikumar@tifr.res.in

‡Indian Institute of Technology, Guwahati. E-mail: saswata.sh@iitg.ernet.in

§Tata Institute of Fundamental Research, Mumbai. E-mail: rakesh@tifr.res.in

every hyperedge $e \in E(H)$ has a vertex colored red and a vertex colored blue, that is, no hyperedge is *monochromatic*. Note that the hypergraph two-coloring problem can be viewed as a constraint satisfaction problem where the clauses are of a specific kind (the Not-All-Equal predicate).

In the special case of graphs (which we may view as two-uniform hypergraphs) two-colorability is easy to characterize and establish: the graph is two-colorable if and only if it does not have an odd cycle; one can find a two-coloring in linear time and $O(|V|)$ space with random access to the input; if the edges are streamed one at a time, a two-coloring can be constructed with $O(|V| \log |V|)$ space and $O((|E| + |V|) \log^2 |V|)$ bit operations.

For the case of hypergraphs, the situation is not as simple. Using the probabilistic method, Erdős [2] showed that any n-uniform hypergraph with fewer than 2^{n-1} hyperedges is two-colorable: a random two-coloring is valid with positive probability; furthermore, this randomized method can be derandomized using the method of conditional probabilities. Erdős [2] later showed that there are n-uniform hypergraphs with $\Theta(n^2)$ vertices and $\Theta(n^2 2^n)$ hyperedges that are *not* two-colorable. Both these bounds remained unchanged for some time, until it was first improved by Beck [1], and further improved by Radhakrishnan and Srinivasan [4], who showed that any hypergraph with fewer than $0.7\sqrt{\frac{n}{\ln n}}2^n$ hyperedges is two-colorable. They also provided a polynomial-time randomized algorithm (and its derandomization) for coloring such hypergraphs. The algorithm has *one-sided* error, i.e., with some small probability δ that can be made arbitrarily small, it declares failure, but otherwise it always outputs a valid two-coloring. Recently, Cherkashin and Kozik [13] showed that a remarkably simple randomized algorithm achieves this bound. Erdős and Lovász [3] have conjectured that any hypergraph with fewer than $n2^n$ hyperedges is two-colorable. In related work, Achlioptas et al. [9] studied when, in terms of the number of hyperedges, does a *randomly* chosen n-uniform hypergraph stop being two-colorable, and relate this to similar questions for random n-SAT. The above results are formally stated in terms of bounds on the number $q(n)$, the minimum number of hyperedges in a *non-two-colorable* hypergraph.

In its original description, the delayed recoloring algorithm of Radhakrishnan and Srinivasan [4] assumes that the entire hypergraph is available in memory, its vertices and hyperedges and may be accessed quickly. Unfortunately, the number of hyperedges can be exponentially larger than the number of vertices. For example, the number of hyperedges can be as high as $\Omega(2^n)$ even when there are just $O(n)$ vertices, and it may be unrealistic to expect the entire hypergraph to be available in the main memory. In this paper, we ask if the performance of the delayed recoloring algorithm can be replicated when the work space of the algorithms is limited to (say) a polynomial in the size of vertex set.

To study this question, we consider the hypergraph two-coloring problem in the streaming model, where the data arrives as a stream or is stored in external memory, and an algorithm with limited local work space analyzes it by making a small number of sequentially passes over it. The resources whose use one tries to minimize in this model are: the number of passes over the data, the amount of local memory required, and the maximum processing time for any data item. On receiving the complete data set, the algorithm must decide as fast as possible either to start another pass or stop with an output. The algorithm could be either deterministic or randomized, and in the latter case, its output needs to be accurate with high probability over the algorithm's internal coin tosses. A number of important algorithms have been developed in the streaming model of computation, e.g., estimating frequency moments [8] and heavy hitters [12] of a data stream. These algorithms find applications in the context of network routing, where large amount of data flows through a router but each of them has limited memory.

Further motivation for our work comes from the *semi-streaming model* for graphs that has recently been considered widely in the literature. Proposed by Muthukrishnan [17], this model looks at solving fundamental problems on graphs such as Max-Cut, $s-t$-connectivity, shortest paths, etc. (see, for e.g. the works [18, 19, 20] and references therein) when the edges are streamed in one at a time. The algorithm only has $O(|V|\text{poly}\log(|V|)$ bits of workspace, and limited number of passes over the edge-stream, which means that it cannot store the entire graph in memory.

Our model for hypergraphs is an extension of this setting, when we have a stream of *hyperedges* coming in. The parameter of interest to us is the *uniformity* n of the hypergraph, and we assume the number of vertices is $\text{poly}(n)$: this is the natural regime in which all the results known for two-coloring hypergraphs are stated. The hyperedges are made available one at a time as a list of n vertices each, where every vertex is represented using B bits. Furthermore, we assume that $\text{poly}(n)$ bits (equivalently, $\text{poly}(|V|)$ bits are available in the algorithm's work space. Notice that the number of hyperedges could now be *exponential* in n (or subexponential in $|V|$), and hence we allow for a space that is $\text{poly}(|V|)$, rather than close to linear in the number of vertices like in the semi-streaming model for graphs.

1.1 Deterministic streaming algorithms:

In the first part of the paper, we investigate the space requirements of *deterministic* streaming algorithms for the hypergraph two-coloring problem. It seems reasonable to conjecture that any deterministic two-coloring one-pass algorithm must essentially store all the hyperedges before it can arrive at a valid coloring. Let $s^D(v, q, n)$ be the minimum space used by a deterministic (one-pass) streaming algorithm that on receiving an n-uniform hypergraph H on v vertices and q hyperedges produces a proper two-coloring of H.

Theorem 1.1. $s^D(n^2, q, n) = \Omega(q/n)$.

Note that when $q \leq 0.7\sqrt{\frac{n}{\ln n}}2^n$, the graph is guaranteed to have a two-coloring. However, the above theorem shows if the number of hyperedges is large, a two-coloring (though it is guaranteed to exist) cannot be found efficiently by a deterministic streaming algoirithm.

Lower bounds for space bounded computations often follow from lower bounds for associated communication complexity problems. The above result is also obtained using this strategy. However, the communication complexity problem turns out to be somewhat subtle; in particular, we are not able to directly reduce it to a well-known problem and refer to an existing lower bound. We conjecture that no deterministic algorithm can do substantially better if it is allowed only a constant number of passes over the input. Note, however, that proving this might be non-trivial, because the corresponding two-round deterministic communication problem has an efficient protocol; see Section 3.1 for details.

1.2 Randomized streaming algorithms:

We show that a version of the delayed recoloring algorithm can be implemented in the streaming model, and provides essentially the same guarantees as the original algorithm. Let $s^R(v, q, n)$ be the minimum space used by a randomized (one-pass) streaming algorithm that on receiving an n-uniform hypergraph H on v vertices and q hyperdges with probability at least $\frac{3}{4}$ (say) produces a proper two-coloring of H (or declares failure). We suppose that each vertex is represented using B bits, and each edge is represented as an n-tuple of vertices.

Theorem 1.2. $s^R(v, \frac{1}{10}\sqrt{\frac{n}{\ln n}}2^n, n) = O(vB)$. *Furthermore, the corresponding randomized algorithm maintains a coloring of the vertices encountered, and updates this coloring in time $O(nvB)$ per hyperedge. If the hypergraph has at most $\frac{1}{10}\sqrt{\frac{n}{\ln n}}2^n$ hyperedges, then with high probability the two-coloring is valid. If the two-coloring being maintained valid, the algorithm declares failure (the algorithm never outputs an invalid coloring).*

1.3 The number of vertices in non-two-colorable hypergraphs

The upper bound for Property B, $q(n) = O(n^2 2^n)$, was shown by Erdős by exhibiting a hypergraph with $O(n^2 2^n)$ hyperedges that is not two-colorable; this hypergraph needed n^2 vertices. We ask if there are such non-two-colorable hypergraphs with $o(n^2)$ vertices.

Theorem 1.3. *(a) Let t be such that $4 \leq t \leq \frac{n^2}{2n-1}$. An n-uniform hypergraph with at most n^2/t vertices and at most $2^{n-1}\exp(t/8)$ hyperedges is two-colorable. The corresponding two-coloring can be found using an efficient randomized streaming algorithm: $s^R\left(\frac{n^2}{t}, 2^{n-1}\exp(t/8), n\right) = O(n^2/t)$.*

(b) If $t \leq \frac{n^2}{2n-1}$, there is a graph with n^2/t vertices and $O(\frac{n^2}{t}\cdot 2^n \exp(\frac{tn}{n-2t}))$ hyperedges that is not two-colorable.

Setting $t = 8\ln(2n)$ in Theorem 1.3(a) above implies that for any n-uniform hypergraph with $v \leq \frac{n^2}{8\ln(2n)}$ vertices and $q < n^2 2^n$ hyperedges, there is a randomized one-pass streaming algorithm that outputs a two-coloring with high probability. This algorithm requires $O(\frac{n^2}{\ln n})$ space at any instant and $O(n)$ processing time after reading each hyperedge. A comparison of this to the result of Erdős [2], which gave a construction of a non-two-colorable hypergraph with the same $q = n^2 2^n$ hyperedges, but on $\Theta(n^2)$ vertices shows that the number of vertices being $\Omega(n^2)$ was crucial, and anything significantly smaller would not have worked. Furthermore, Theorem 1.3 (b) recovers Erdős' bound when $t = O(1)$, and generalizes it to make explicit the dependence of the Erdős' upper bound on $q(n)$ on the vertex-set size.

Organisation of the rest of the paper: In Section 2 we introduce the notation. In Section 3, we establish Theorem 1.1, showing the limitation of deterministic streaming algorithms. In Section 4, we prove Theorem 1.2 by showing how a version of the off-line delayed recoloring algorithm of [4] (see Section 4.1) can, in fact, be implemented efficiently in the streaming model. In Section 5, we prove the results in Theorem 1.3 regarding non-two-colorability. We conclude with some remarks and open problems.

2 Notation

Our n-uniform hypergraphs will be denoted by (V, E), where V is the set of vertices, and $E \subseteq \binom{V}{n}$ is the set of hyperedges of H. For a hypergraph $H = (V, E)$, we use v for $|V|$, and q for $|E|$. In our setting, v will typically be a small polynomial in n and q will be exponential in n. For any $k \in \mathbb{N}$, we use the notation $[k] \coloneqq \{1, \ldots, k\}$.

A valid *two-coloring* of the hypergraph H is an assignment $\chi : V \to \{\text{Red}, \text{Blue}\}$ that leaves no edge monochromatic, i.e., $\forall e \in E$, $\exists i, j \in e$ such that $\chi(i) \neq \chi(j)$. A hypergraph that admits a valid two-coloring is said to be *two-colorable,* or equivalently, to have *Property B.*

We will be interested in space-efficient streaming algorithms for finding two-colorings of hypergraphs. Consider a hypergraph with edge set $E = \{e_1, \ldots, e_q\}$. The hyperedges are made available to the algorithm one at a time in some order $e_{i_1}, \ldots, e_{i_q}$. To keep the problem general, we will not assume that the vertex set is fixed in advance. The algorithm will become aware of the vertices as they arrive as part of the stream of hyperedges. We will assume that each vertex is encoded using B bits. The goal is to design a space-efficient algorithm (deterministic or randomized) that can output a valid coloring for the entire hypergraph once all the hyperedges have *passed*. We may allow multiple passes over the input data. We call the algorithm a r-pass streaming algorithm, if it outputs a valid coloring after making r passes over the input stream. By space-efficient, we mean that the algorithm uses poly(n) internal workspace. We will assume that n is large, say at least 100.

Denote by $s_r^D(v, q, n)$ the minimum space used by a deterministic r-pass streaming algorithm that on receiving an n-uniform hypergraph H on v vertices and q hyperedges produces a proper two-coloring of H. Similarly, $s_r^R(v, q, n)$ is the minimum space used by a randomized r-pass streaming algorithm that on receiving an n-uniform hypergraph H on v vertices and q hyperedges with probability at least $\frac{3}{4}$(say) produces a proper two-coloring of H (or declares failure). When r is omitted from the subscript, it is assumed $r = 1$.

3 Deterministic streaming algorithms

We first show lower bounds in the deterministic setting. We recall the routine translation of an efficient streaming algorithm to a communication complexity protocol [8], with a view to proving lower bounds. There are two computationally unbounded players Alice and Bob, who both know of a relation $R \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$. Alice receives an input $x \in \mathcal{X}$ and Bob gets $y \in \mathcal{Y}$; in the beginning, neither player is aware of the other's input. Their goal is to exchange bits according to a fixed *protocol* and find a $z \in \mathcal{Z}$ so that $(x, y, z) \in R$. The communication complexity of R is the minimum number of bits that Alice and Bob exchange in any valid protocol for the worst case input pair (x, y). Several generalizations of this model can be defined with $k >= 3$ players, we will define our specific model below. For more details on communication complexity in general, please consult the book by Kushilevitz and Nisan [7].

To show lower bounds for hypergraph two-coloring, we define the class of communication problems $\mathcal{H}(v, q, k)$ [1].

Definition 3.1 (Problem class $\mathcal{H}(v, q, k)$)**.** For $k >= 2$, an instance $I \in \mathcal{H}(v, q, k)$ has k players $P_1, \ldots, P_k$. Each player P_i has a subset $E_i \subseteq E$ of some hypergraph $H = (V, E = E_1 \cup E_2 \cup \cdots \cup E_k)$, with $|V| = v$ and $|E_i| \leq q$. The communication is done in sequential order: starting with P_1, P_i sends a message to P_{i+1} for $i \in \{1, \ldots, k-1\}$. This sequence of communication constitutes a round. In a multiple-round protocol, P_k may start a new round by sending back a message to P_1, who would continue communication in the above order. In a valid r-round protocol Π, some P_i in the course of the r-th round will output a coloring χ that is valid for H. Denote by $\Pi(I, P_i, l)$ the communication sent by P_i in round l on instance I. We define the r-round communication complexity of $\mathcal{H}(v, q, k)$ as follows:

$$\mathrm{CC}_r(\mathcal{H}(v, q, k)) = \min_{\substack{\Pi:\ r-\text{round} \\ \text{valid protocol}}} \max_{\substack{I \in \mathcal{H}(v,q,k), \\ i \in [k], l \in [r]}} |\Pi(I, P_i, l)|$$

[1] Since all hypergraphs we consider are n-uniform, we will not explicitly state n as a parameter

Remark 3.2. In this definition, we consider the maximum communication by any single player instead of total communication because this quantity is related more closely to the space requirement of streaming protcols (see Proposition 3.3 below). If the maximum communication by any player is s bits, then the total communication is bounded above by $kr \cdot s$ bits for an r round protocol. Besides, for the values of k and r that we consider, this $O(kr)$ blowup is immaterial.

Note that when $k = 2$ and $r = 1$, we get the two-player one-round model (where we traditionally call P_1 as Alice and P_2 as Bob): Alice sends a message m to Bob depending on her input and Bob outputs a coloring looking at m and his input. Our lower bound for hypergraph coloring in the streaming model will rely on the following well-known connection between streaming and communication complexity introduced in [8].

Proposition 3.3 ([8])**.** *For any $k \in \mathbb{Z}$, $k \geq 2$, we have:*

$$\mathrm{CC}_r(\mathcal{H}(v, q, k)) \leq s_r^D(v, q, n)$$

Thus, to establish Theorem 1.1, it is enough to show an appropriate lower bound on $\mathrm{CC}_1(\mathcal{H}(v, q, k))$. We start with the two-player case, which already introduces most of the ideas.

Theorem 3.4. $\mathrm{CC}_1(\mathcal{H}(n^2, q, 2)) = \Omega(\frac{q^2}{2^n n^4})$.

Remark 3.5. Note that the above theorem gives a non-trivial lower bound only when $q \gg n^2 2^{n/2}$. Ideally, we would expect a lower bound that is linear in q, for all values of q.

Proof. Consider a valid one-round protocol for the two-player problem where Alice sends a message m from a set $\mathcal{M}$ of possible messages. Let the input to the protocol be (H_A, H_B), where H_A and H_B are hypergraphs, each with q hyperedges on a common vertex set $[n^2]$. For every hypergraph H_B that Bob receives, he must output a coloring $\chi = f(m, H_B)$ based on some deterministic function f. For $m \in \mathcal{M}$, define

$$L(m) = \{f(m, H_B) : H_B \text{ is an input for Bob}\}.$$

It is easy to see that Bob may identify the message m with the list $L(m)$; on receiving m, he must find a proper coloring for H_B from $L(m)$. Thus, for every $m \in \mathcal{M}$, we have the following.

Completeness for Bob: Every possible input hypergraph to Bob (i.e. all hypergraphs on q hyperedges) should have a valid coloring in $L(m)$.

Soundness for Alice: Let $P(m) = \{H_A : \text{Alice sends the message } m \text{ for input } H_A\}$. Then every $\chi \in L(m)$ should be valid for every hypergraph $H_A \in P(m)$.

We will show that these two conditions imply the claimed lower bound on $|\mathcal{M}|$.

Definition 3.6 (Shadows)**.** Given a coloring χ, we define its *shadow* $\Delta(\chi)$ to be the set of all possible hyperedges that are monochromatic under χ. The shadow of a list L of colorings is

$$\Delta(L) = \bigcup_{\chi \in L} \Delta(\chi)$$

Note that in the above definition, the shadow $\Delta(\chi)$ collects *all possible* monochromatic hyperedges under χ, so it depends only on the coloring χ, and not on any hypergraph. Similarly, $\Delta(L)$ is also a collection of hyperedges and does not depend on any hypergraph; in particular, if a hypergraph H_A is monochromatic under every coloring in L, then none of H_A's hyperedges can appear in $\Delta(L)$.

In the following, assume that n is large.

Claim 3.7. *For every coloring χ, we have*

$$\Delta(\chi) \geq \frac{1}{10 \cdot 2^n} \binom{n^2}{n}. \tag{3.1}$$

Proof of Claim. : One of the two color classes χ has at least $\lceil n^2/2 \rceil$ vertices. It follows that

$$\Delta(\chi) \geq \binom{\lceil n^2/2 \rceil}{n} \geq \frac{1}{10 \cdot 2^n} \binom{n^2}{n}.$$

□

We next observe that the *completeness condition* for Bob imposes a lower bound on $\Delta(L(m))$.

Claim 3.8. *For every $m \in \mathcal{M}$,*

$$|\Delta(L(m))| \geq \frac{q}{10n^2 2^n} \binom{n^2}{n}. \tag{3.2}$$

Proof of Claim. : Suppose the claim does not hold, that is,

$$|\Delta(L(m))| < \frac{q}{10n^2 2^n} \binom{n^2}{n}. \tag{3.3}$$

Choose a random hypergraph H_B by choosing q hyperedges randomly from $\Delta(L(m))$. We will say that the hypergraph *hits* $\chi \in L(m)$, if at least one of its hyperedges falls in $\Delta(\chi)$, otherwise we say it *misses* χ. For each $\chi \in L(m)$, we have, using the bounds (3.1) and (3.3), that

$$\begin{aligned} \Pr_H[H \text{ misses } \chi] &\leq \left(1 - \frac{|\Delta(\chi)|}{|\Delta(L(m))|}\right)^q \\ &\leq \left(1 - \frac{n^2}{q}\right)^q \\ &\leq \exp(-n^2) \end{aligned}$$

Since there are at most 2^{n^2} colorings, the union bound yields:

$$\Pr_H[\exists \chi \in L(m) : H \text{ misses } \chi] \leq 2^{n^2} \exp(-10n^2) \ll 1$$

Thus there exists a hypergraph H with q hyperedges that hits every coloring $\chi \in L(m)$, that is every coloring in $L(m)$ is invalid for H. This, however, violates Completeness for Bob, proving the Claim. □

We can now complete the proof of the theorem. Consider a random hypergraph H for Alice, obtained by choosing each of its hyperedges uniformly at random from the set of all hyperedges. Since Alice sends some $m \in \mathcal{M}$ for *every* hypergraph, the *soundness condition* for Alice implies that

$$1 = \Pr_H[\exists m : H \text{ misses all } \chi \in L(m)] \leq |\mathcal{M}|(1 - \frac{q}{10n^2 2^n})^q \leq |\mathcal{M}| \exp(-\frac{q^2}{10n^2 2^n}),$$

where we used Claim 3.8 to justify the first inequality.

Taking logarithms on both sides yields the desired lower bound $\log |\mathcal{M}| = \Omega(\frac{q^2}{n^4 2^n})$. □

As remarked earlier, the above communication lower bound implies that the space required by a deterministic streaming algorithm to find a valid coloring is exponential in n even for hypergraphs that have very simple randomized coloring strategies; it does not yield any such lower bound for hypergraphs that have fewer than $2^{n/2}$ hyperedges. In order to overcome this limitation, we generalize the analysis above to the k-player hypergraph coloring problem. We will see later that for $q \leq 2^{n/2}$ there do exist efficient two-player one-round protocols, so we could not have proved our lower bounds while restricting our attention to the two-player setting.

Proof of Theorem 1.1:

Theorem 1.1 will follow from the following theorem:

Theorem 3.9. *Let $k \geq 1$.*

$$\mathrm{CC}_1(\mathcal{H}(v, q, k+1)) = \Omega\left(q \left[\frac{q}{v} \frac{\binom{v/2}{n}}{\binom{v}{n}}\right]^{\frac{1}{k}}\right)$$

Proof. Recall that the protocol has $k+1$ players, $P_1, P_2, \ldots, P_k, P_{k+1}$. player P_i receives a hypergraph H_i with q hyperedges over the vertex set $[v]$. The communication starts with P_1, who sends a message m_1 of length ℓ_1 to player P_2; in the i-th step, P_i sends a message of length ℓ_i to P_{i+1}. In the end, P_{k+1} produces a coloring for the hypergraph $H_1 \cup H_2 \cup \cdots \cup H_{k+1}$. It will be convenient to view this coloring as a message sent by player P_{k+1}, and set $\ell_{k+1} = v$ (the number of bits needed to describe a coloring).

In a $(k+1)$-player protocol, after the messages sent by the first i players have been fixed, we have a list of colorings that may still be output at the end; we use the following notation to refer to this list (recall that m_{k+1} is a coloring):

$$L(m_1, m_2, \ldots, m_i) = \{m_{k+1} : \text{ for some input } H_1, H_2, \ldots, H_{k+1} \text{ the transcript is of the form } (m_1, m_2, \ldots, m_i, \ldots, m_k, m_{k+1})\}.$$

In particular, by considering the situation at the beginning of the protocol (when no messages have yet been generated), we let $L_0 = \{\chi : m_{k+1} = \chi \text{ is output by the protocol on some input}\}$. Let $s_0 = |\Delta(L_0)|/\binom{v}{n}$, and for $i = 1, \ldots, k+1$, let

$$s_i(m_1, \ldots, m_i) = |\Delta(L(m_1, m_2, \ldots, m_i))| \Big/ \binom{v}{n}$$

and

$$s_i = \min_{(m_1,m_2,\ldots,m_i)} s_i(m_1,\ldots,m_i).$$

Here the minimum is taken over all possible sequences of first i messages that arise in the protocol. In particular, s_0 corresponds to the union of shadows of all colorings ever output by the protocol, and s_{k+1} corresponds to the shadow of the output corresponding to the transcript (that is, the shadow of the last message, which is a coloring).

Claim 3.10.

$$s_0 \leq 1 \tag{3.4}$$

$$\forall i \in \{0,\ldots,k\}: \quad s_i \geq s_{i+1}\frac{q}{\ell_{i+1}} \tag{3.5}$$

$$s_{k+1} \geq \frac{\binom{\lceil v/2 \rceil}{n}}{\binom{v}{n}}. \tag{3.6}$$

Proof. Inequality (3.4) is immediate from the definition of s_0.

For Claim (3.5), we use ideas similar to those used in the proof of Claim 3.8 for two-player protocols. Fix $m_1, m_2, \ldots, m_i$. We will show that $s_i(m_1, m_2, \ldots, m_i)$ is at least the right hand side of (3.5). Pick a random hypergraph H (which we will consider as a possible input to P_{i+1}) as follows: choose q hyperedges independently and uniformly from the set $\Delta(L(m_1,\ldots,m_i))$. When H is presented to P_{i+1}, it must respond with a message m_{i+1}. None of the colorings that P_{k+1} produces after that can include any edge of H in its shadow, that is, $\Delta(L(m_1,\ldots,m_{i+1})) \cap H = \emptyset$. For each valid choice m for m_{i+1}, we have $\Delta(L(m_1,\ldots,m_i,m)) \subseteq \Delta(L(m_1,\ldots,m_i))$ and

$$\Pr[H \cap \Delta(L(m_1,\ldots,m_i,m)) = \emptyset] \leq \left(1 - \frac{s_{i+1}(m_1,\ldots,m_i,m)}{s_i(m_1,\ldots,m_i))}\right)^q \leq \left(1 - \frac{s_{i+1}}{s_i(m_1,\ldots,m_i))}\right)^q.$$

Thus,

$$1 = \Pr_H[\exists m_{i+1} : \Delta(L(m_1,\ldots,m_{i+1})) \cap H = \emptyset] \leq 2^{\ell_{i+1}}\left(1 - \frac{s_{i+1}}{s_i(m_1,\ldots,m_i))}\right)^q$$

This yields $\exp(\ell_{i+1} - \frac{s_{i+1}q}{s_i(m_1,\ldots,m_i)}) \geq 1$, giving $s_i(m_1,\ldots,m_i) \geq s_{i+1}\left(\frac{q}{\ell_{i+1}}\right)$. By minimizing over valid sequences $(m_1,\ldots,m_i)$, we justify our claim.

Claim 3.6 follows from the fact the shadow of every coloring has at least $\binom{\lceil v/2 \rceil}{n}$ hyperedges. □

By combining parts (3.4) and (3.5), we obtain

$$1 \geq q^{k+1}s_{k+1} \Big/ \prod_{i=1}^{k+1} \ell_i \,.$$

The theorem follows from this by using (3.6), noting that $\ell_{k+1} \leq v$, and for $i = 1, 2, \ldots, k$: $\ell_i \leq \max_{i\in[k]} \ell_i$. □

Corollary 3.11 (Restatement of Theorem 1.1)**.** *Every one-pass deterministic streaming algorithm to two-color an n-uniform hypergraph with at most q hyperedges requires $\Omega(\frac{q}{n})$ bits of space.*

Proof. Setting $k = n$ in the previous result immediately yields that for hypergraphs on $v = n^2$ vertices and at most $q = (n+1)q'$ hyperedges, the communication required is $\Omega(q')$. The lower bound for streaming algorithms then follows from Proposition 3.3. □

3.1 Deterministic communication protocols for the two-coloring problem

In the previous section, we derived our lower bound for the deterministic streaming algorithms by invoking *multi-player* communication complexity, because the two-player lower bound did not give us a non-trivial lower bound for hypergraphs with fewer than $2^{n/2}$ hyperedges. In this section, we first show an upper bound in the two-player setting, which shows that it was essential to consider the multi-player setting in order to get the stronger lower bound. Next, we consider two-round protocols, for they are related to two-pass streaming algorithms. We show below, perhaps surprisingly, that the two-round two-player deterministic communication complexity for the problem is poly(n). However, we do not have a streaming algorithm with a matching performance.

Theorem 3.12. $\mathrm{CC}_1(\mathcal{H}(v, 2^{n/2}, 2)) = \mathrm{poly}(n)$

Proof. Alice and Bob will base their protocol on a special collection of lists of colorings $\mathcal{L} = \{L_1, \ldots, L_r\}$, with $r = 2^n$. Suppose Alice's hypergraph is H_A and Bob's hypergraph is H_B. The protocol will have the following form.

Alice Alice sends an index i of a list $L_i \in \mathcal{L}$ such that every coloring in L_i is valid for H_A.

Bob Bob outputs a coloring $\chi \in L_i$ that is valid for H_B.

We next identify some properties on the collection of lists that easily imply that the protocol above produces a valid two-coloring.

Definition 3.13 (Good Lists)**.** (a) A collection of lists $\mathcal{L}$ is good for Bob, if for every hypergraph H_B, in every list L in $\mathcal{L}$ there is a valid coloring for H_B in L.

(b) A collection of lists $\mathcal{L}$ is good for Alice, if for every hypergraph H_A, there is some list L in $\mathcal{L}$, such that every coloring in L is valid for H_A.

Note that good collections are defined differently for Alice and Bob. With this definition, the existence of a collection of lists that is good for both Alice and Bob would furnish a one-way communication protocol with $\lceil \log r \rceil$ bits of communication, and establish the theorem. Such a *good* collection is shown to exist in Lemma 3.14 below. □

Lemma 3.14. *Suppose H_A and H_B are restricted to have at most $2^{n/2}$ hyperedges. Then, there exists a collection of $\mathcal{L} = \{L_1, \ldots, L_r\}$ of $r = 2^n$ lists, each with $k = \lceil 62^{n/2} \log v \rceil$ colorings, which is good for Alice and Bob.*

Proof. We pick $\mathcal{L}$ by picking r lists randomly: list L_i will be of the form $\{\chi_{i1}, \ldots, \chi_{ik}\}$, where each coloring is chosen independently and randomly from the set of all colorings. We will show that with positive probability $\mathcal{L}$ is a good collection of lists. We will separately bound the probability that $\mathcal{L}$ fails to satisfy conditions (a) and (b) above. We may restrict attention to hypergraphs that have exactly $q = \lfloor 2^{n/2} \rfloor$ hyperedges.

(a) We first bound the probability that $\mathcal{L}$ is not good for Bob. Consider one list $L_i \in \mathcal{L}$.

$$\begin{aligned} \Pr[L_i \text{ is not good for Bob }] &= \Pr[\exists H_B \text{ such that none of the } k \text{ colorings are valid for } H_B] \\ &\leq \binom{\binom{v}{n}}{q} \left(\frac{q}{2^{n-1}}\right)^k. \end{aligned}$$

Since $\mathcal{L}$ is a collection of r lists, we have

$$\begin{aligned}\Pr[\mathcal{L} \text{ is not good for Bob}] &= \Pr[\exists L_i \in \mathcal{L} \text{ is not good for Bob}] \\ &\leq r\binom{\binom{v}{n}}{q}\left(\frac{q}{2^{n-1}}\right)^k \qquad (3.7) \\ &\leq 2^{(nq\log v+\log r)-k(n-\log q-1)}. \qquad (3.8)\end{aligned}$$

To make the right hand side of (3.8) small, we will later choose k such that

$$k \gg \frac{nq\log v+\log r}{n-\log q-1}. \qquad (3.9)$$

(b) We next bound the probability that $\mathcal{L}$ is not good for Alice. Fix a hypergraph H_A for Alice, and consider $L_i \in \mathcal{L}$.

$$\begin{aligned}\Pr_{L_i}[\text{every } \chi \in L_i \text{ is valid for } H_A] &\geq \left(1-\frac{q}{2^{n-1}}\right)^k \\ &\geq \exp\left(-\frac{qk}{2^{n-1}-q}\right). \quad (\text{using } 1-x \geq \exp\left(-\tfrac{x}{1-x}\right))\end{aligned}$$

Since $\mathcal{L}$ is a collection of r such lists chosen independently and there are $\binom{\binom{v}{n}}{q}$ choices for H_A, we have

$$\begin{aligned}\Pr[\mathcal{L} \text{ is not good for Alice}] &\leq \binom{\binom{v}{n}}{q}\left(1-\exp\left(-\frac{qk}{2^{n-1}-q}\right)\right)^r \\ &\leq \exp\left(qn\ln v - r\exp\left(-\frac{qk}{2^{n-1}-q}\right)\right). \qquad (3.10)\end{aligned}$$

To make the right hand side of (3.10) small, we will choose r such that

$$r \gg (qn\ln v)\exp\left(\frac{qk}{2^{n-1}-q}\right). \qquad (3.11)$$

Now, suppose $v = \text{poly}(n)$. Then one can verify that if we set $k = \lceil 6q\log v\rceil$ and $r = 2^n$, then (3.9) and (3.11) both hold for all large n. It follows that the required collection of lists exists. □

3.1.1 Circuit upper bounds and a 2-round communication protocol

Next we present our efficient two-player, two-round communication protocol for hypergraph two-coloring.

Theorem 3.15. $\text{CC}_2(\mathcal{H}(v, 2^n/8, 2)) = \text{poly}(n)$

We will prove Theorem 3.15 by exploiting the connection between circuit complexity and *Karchmer-Wigderson games* [14].

Definition 3.16 (Karchmer-Wigderson game)**.** Given a monotone boolean function $f : \{0,1\}^N \to \{0,1\}$, the Karchmer-Wigderson communication game G_f between two players Alice and Bob is the following: Alice gets an input $x \in f^{-1}(0)$, and Bob gets an input $y \in f^{-1}(1)$, and they both know f. The goal is to communicate and find an index $i \in [N]$ such that $x_i = 0$ and $y_i = 1$ (such an index exists since f is monotone). The communication complexity of this game is denoted by $\text{CC}(G_f)$ (or specifically $\text{CC}_r(G_f)$ for an optimal r-round protocol).

Karchmer and Wigderson make the following connection:

Theorem 3.17 ([14]). *If a monotone boolean function f has a depth-d, size s circuit, then* $\mathrm{CC}_{d-1}(G_f) = O(d \log s)$.

The proof of Theorem 3.15 will be based on the circuit complexity of *Approximate Majority functions*. Define the partial function $\mathrm{ApprMaj}_N$ on a subset of $\{0,1\}^N$ as follows: $\mathrm{ApprMaj}_N(x) = 1$ when $|x| \geq 2N/3$, and $\mathrm{ApprMaj}_N(x) = 0$ when $|x| \leq N/3$. A function $f : \{0,1\}^N \to \{0,1\}$ is an Approximate Majority function if it satisfies: $f(x) = 1$ when $|x| \geq 2N/3$, and $f(x) = 0$ when $|x| \leq N/3$. We need the following Lemma.

Lemma 3.18. *For any Approximate Majority function $f : \{0,1\}^{2^v} \to \{0,1\}$, we have:*

$$\mathrm{CC}_r\left(\mathcal{H}(v,\ 2^n/8,\ 2)\right) \leq \mathrm{CC}_r(G_f)$$

Proof. We will prove this by reducing $\mathcal{H}(v,\ 2^n/8,\ 2)$ to G_f, for any Approximate Majority function f. For convenience, let $N \coloneqq 2^v$. Suppose Alice receives $H_A = (V, E_A)$, and Bob receives $H_B = (V, E_B)$. Let $L_A \in \{0,1\}^N$ be an indicator vector for Alice that marks which of the N possible two-colorings of V (arranged in some canonical order) are valid for H_A, and let L_B be the corresponding indicator vector for Bob. Since we know that both H_A, H_B have $q \leq 2^n/8$ hyperedges, a randomly chosen coloring will color either hypergraph with probability $\geq 3/4$, which implies $|L_A| \geq 3N/4$ and $|L_B| \geq 3N/4$. To find a coloring valid for both H_A and H_B, Alice and Bob just have to find an index $i \in [2^v]$ such that $L_A(i) = L_B(i) = 1$.

Alice now complements her input to consider $\overline{L_A}$. Since $|\overline{L_A}| \leq N/4$, $\mathrm{ApprMaj}_N(\overline{L_A}) = 0$, whereas $\mathrm{ApprMaj}_N(L_B) = 1$. Clearly, playing the Karchmer-Wigderson game for any Approximate Majority function f on N bits, $\overline{L_A}, L_B$ would find a index i where $L_A = L_B = 1$, and consequently yield a valid coloring for $H_A \cup H_B$. □

Hence, a small depth-3 monotone circuit that computes $\mathrm{ApprMaj}_{2^v}$ on its domain would yield an efficient 2-round protocol for $\mathcal{H}(v,\ 2^n/8,\ 2)$, using Theorem 3.17. Ajtai [15] showed that such a circuit exists, and Viola [16] further showed uniform constructions of such circuits.

Theorem 3.19. *[16] There exist monotone, uniform* $\mathrm{poly}(N)$*-sized depth-3 circuits for* $\mathrm{ApprMaj}_N$ *on N input bits.*

Proof (of Theorem 3.15). The proof of Theorem 3.15 is now immediate from Theorems 3.17, 3.19 and Lemma 3.18, since v is polynomial in n. □

4 Streaming algorithms for hypergraph two-coloring

In this section, we present streaming algorithms that come close to the performance of the randomized off-line algorithm of Radhakrishnan and Srinivasan [4]. We first point out why this algorithm, as stated, cannot be implemented with limited memory. Next, we show how an alternative version can be implemented using a small amount of memory. This modified version, however, does return colorings that are not valid (though it does this with small probability); we show how by maintaining a small amount of additional information, we can derive an algorithm that returns and valid coloring with high probability or returns failure, but never returns an invalid coloring.

In order to describe our algorithms, it will be convenient to use u, w to denote vertices in the hypergraph. Also, in this section, we will explicitly use $|V|$ to denote the size of the vertex set V,

and $|E|$ to denote the number of hyperedges in the hypergraph. The vertex set of the hypergraph will be identified with $[|V|]$, and B will denote the number of bits required to represent a single vertex in V.

4.1 The delayed recoloring algorithm

Radhakrishnan and Srinivasan [4] showed that by introducing delays in the recoloring step of an algorithm originally proposed by Beck [1] one can two-color hypergraphs with more hyperedges than it was possible before.

Theorem 4.1. *Let $H = (V, E)$ be an n-uniform hypergraph with at most $1/10\sqrt{n/\ln n} \times 2^n$ hyperedges. Then H is two-colorable; also a proper two-coloring can be found with high probability in time $O(\mathrm{poly}(|V| + |E|))$.*

Algorithm 1 (Off-line delayed recoloring algorithm)

1: Input $H = (V, E)$.
2: For all $u \in V$, independently set $\chi_0(u)$ to Red or Blue with probability $\frac{1}{2}$.
3: Let $\mathcal{M}_0$ be the set of hyperedges of H, that are monochromatic under χ_0.
4: For all $u \in V$, independently set $b(u)$ to be 1 with probability p, and 0 with probability $1 - p$.
5: Let π be one of the $|V|!$ permutations of V, chosen uniformly at random.
6: **for** $i = 1, 2, \ldots, |V|$ **do**
7: $\quad\chi_i$ is obtained from χ_{i-1} by retaining the colors of all vertices except perhaps $\pi(i)$. If $b(\pi(i)) = 1$, and some hyperedge containing $\pi(i)$ was monochromatic under χ_0 and all its vertices still have the same color in χ_{i-1}, then $\chi_{i-1}(\pi(i))$ is flipped to obtain $\chi_i(\pi(i))$.
8: **end for**
9: Output $\chi_f = \chi_{|V|}$.

The term *off-line* refers to the fact that we expect all hyperedges to remain accessible throughout the algorithm. The term *delayed recoloring* refers to the fact that vertices are considered one after another, and are recolored only if the initially monochromatic edge they belonged to has not been set right by a vertex recolored earlier.

For a coloring χ, let $\mathcal{B}(\chi)$ be the set of hyperedges that are colored entirely blue under χ, and let $\mathcal{R}(\chi)$ be the set of hyperedges colored entirely red under χ. In [4], the following claims are established for this algorithm under the assumption that $|E| \le \frac{1}{10}\sqrt{\frac{n}{\ln n}}2^n$. In the following, the parameter p denotes the probability that the bit $b(v)$ is set to 1 in Algorithm 1.

$$\Pr[(\mathcal{B}(\chi_f) \cap \mathcal{B}(\chi_0) \neq \emptyset) \vee (\mathcal{R}(\chi_f) \cap \mathcal{R}(\chi_0) \neq \emptyset)] \le \frac{2}{10}\sqrt{\frac{n}{\ln n}}(1-p)^n \tag{4.1}$$

$$\Pr[\mathcal{B}(\chi_f) \setminus \mathcal{B}(\chi_0) \neq \emptyset] \le \frac{2np}{100 \ln n}; \tag{4.2}$$

$$\Pr[\mathcal{R}(\chi_f) \setminus \mathcal{R}(\chi_0) \neq \emptyset] \le \frac{2np}{100 \ln n}. \tag{4.3}$$

The first inequality (4.1) helps bound the probability that an initially monochromatic hyperedge does not change our attempts at recoloring; the second inequality (4.2) considers the event where an initially non-blue hyperedge become blue because of recoloring; the third inequality (4.3) similarly refers to the event where an initially non-red hyperedge becomes red because of recoloring. These

events together cover all situations where χ_f turns out to be invalid for H. Now, set $p = \frac{1}{2n}(\ln n - \ln\ln n)$ so that we have

$$\frac{2}{10}\sqrt{\frac{n}{\ln n}}(1-p)^n + \frac{4np}{100\ln n} \leq \frac{1}{5} + \frac{1}{50} \leq \frac{11}{50}. \tag{4.4}$$

Thus, the above algorithm produces a valid two-coloring with probability at least $\frac{1}{2}$. Furthermore, since the entire graph is available in this off-line version, we can efficiently verify that the final coloring is valid. By repeating the algorithm $\lceil \log(1/\delta) \rceil$ times, we can ensure that the algorithm produces a valid coloring with probability at least $1-\delta$, and declares failure otherwise (but never produces an invalid coloring).

Implementation in the streaming model: Now, suppose the hyperedges arrive one at a time, each hyperedge is a sequence of n vertices, each represented using B bits. In time $O(Bn)$ per edge, we may extract all the vertices. The corresponding colors and bits for recoloring can be generated in constant time per vertex. Generating the random permutation π is a routine matter, we maintain a random permutation of the vertices received so far by inserting each new vertex at a uniformly chosen position in the current permutation; using specialized data structure, the total time taken for generating π is $O(|V||B|\log|V|)$. Overall, the the algorithms can be implemented in time $\tilde{O}(nB|V||E|)$ (the $\tilde{O}$ hides some polylog($|V|$) factors). The space required is rather large because we explicitly store the hyperedges. Observe, however, that we need store only the hyperedges that are monochromatic after the initial random coloring. Thus, the algorithm can be implemented using $\tilde{O}(|V|B + nB\sqrt{\frac{n}{\ln n}})$ bits of space on the average. The space and time requirements may then be considered acceptable, but this implementation does not allow us to determine if the coloring produced at the end is valid; moreover, it is not clear how we may reduce the error probability by repetition. In the next section, we give an implementation that does not suffer from this deficiency.

4.2 An efficient streaming algorithm

In this section, we modify the randomized streaming algorithm of the previous section so that it uses $O(B|V|)$ bits of space and takes time comparable to the off-line algorithm. This version maintains a running list of vertices as suggested above, but it assigns them colors immediately, and with the arrival of each edge decides if some of its vertices must be recolored. The algorithm is derived from the off-line algorithm. It maintains enough information to ensure that the actions of this streaming algorithm can somehow be placed in one-to-one correspondence with the actions of the off-line algorithm.

We wish to show that Algorithm 2 succeeds in constructing a valid coloring with probability at least $\frac{1}{2}$. To justify this, we compare the actions of this algorithm and the off-line delayed recoloring algorithm stated earlier, and observe that their outputs have the same distribution.

Recall that χ_0 was the initial coloring generated in the Algorithm 1. Let $\hat{\chi}_0$ be the corresponding coloring for Algorithm 2: that is, let $\hat{\chi}_0(u)$ be the color u was first assigned when the first hyperedge containing u appeared in the input. Similarly, let $\hat{b}$ be the sequence of bits generated by the above algorithm, and $\hat{\pi}$ be the final permutation of vertices that results. Notice (χ_0, b, π) and $(\hat{\chi}_0, \hat{b}, \hat{\pi})$ have the same distribution.

In Algorithm 1, once χ_0, b and π are fixed, the remaining actions are deterministic. That is, $\chi_f = \chi_f(\chi_0, b, \pi)$ is a function of (χ_0, b, π). Similarly, once we fix (i.e., condition on) $(\hat{\chi}_0, \hat{b}, \hat{\pi})$, the final coloring $\hat{\chi}_f = \hat{\chi}_f(\hat{\chi}_0, \hat{b}, \hat{\pi})$ is fixed.

Algorithm 2 Delayed recoloring with limited space

Input: The hypergraph $H = (V, E)$ as a sequence of hyperedges $h_1, h_2, \ldots$.

The algorithm will maintain for each vertex u three pieces of information:

(i) its initial color $\chi_0(u)$;

(ii) its current color $\chi(u)$;

(iii) a bit $b(u)$ that is set to 1 with probability p;

Furthermore, it maintains a random permutation π of the vertices received so far.

for $i = 1, 2, \ldots, |E|$ **do**
  Read h_i.
  for $u \in h_i$ **do**
    if u has not been encountered before **then**
      Set $\chi_0(u) = \chi(u)$ to Red or Blue with probability $\frac{1}{2}$. Set $b(u)$ to 1 with probability p and 0 with probability $1 - p$. Insert u at a random position in π.
    end if
  end for
  if h_i was monochromatic under χ_0 and all its vertices have the same color under χ **then**
    Let u be the first vertex (according to the current π) such that $b(u) = 1$.
    if $\chi_0(u) = \chi(u)$ **then**
      Flip $\chi(u)$.
    end if
  end if
end for
Output $\chi_f = \chi_{|V|}$.

Lemma 4.2. *For all χ, b and π, we have $\chi_f(\chi_0, b, \pi) = \hat{\chi}_f(\chi_0, b, \pi)$.*

Proof. Suppose $\chi_f(u) \neq \hat{\chi}_f(u)$ for some vertex u. One of them must equal $\chi_0(u)$. We have two cases.

Suppose $\chi_0(u) = \hat{\chi}_f(u) \neq \chi_f(u)$: Let us also assume that h is a hyperedge that necessitated u's recoloring in the off-line algorithm. This implies that $b(u) = 1$, and $b(w) = 0$ for the vertices w of h that appeared before u in π. Now in the streaming algorithm above, when h is considered, we would find that h is monochromatic in χ_0, and flip $\chi(u)$ unless it is already flipped. Once flipped, the color of u will not change again. Thus, $\hat{\chi}_f(u) \neq \chi_0(u)$—a contradiction.

Suppose $\chi_0(u) = \chi_f(u) \neq \hat{\chi}_f(u)$: Let h be the hyperedge that necessitated the recoloring of u in the above streaming algorithm. This implies that $b(u) = 1$, and all vertices w of h that appear before u in permutation π have $b(w) = 0$. But in such a situation, the original off-line algorithm will find h monochromatic when u is considered, and will flip its color. Thus, $\chi_0(u) \neq \chi_f(u)$—a contradiction.

□

We conclude that (4.4) applies to Algorithm 2 as well.

Corollary 4.3 (Compare Theorem 4.1)**.** *Let $H = (V, E)$ be an n-uniform hypergraph with at most $1/10\sqrt{n/\ln n} \times 2^n$ hyperedges. Then, with probability at least $\frac{1}{2}$ Algorithm 2 produces a valid two-coloring for H.*

Note, however, Algorithm 2 lacks the desirable property that it outputs only valid colorings. A straight-forward check does not seem possible using $O(|V|B)$ RAM space. We show below that with carefully storing some vertices of a few hyperedges, we can in fact achieve this using only $O(nB)$ RAM space.

4.3 An efficient algorithm that produces only valid colorings

As remarked before the events considered in inequalities (4.1), (4.2) and (4.3) cover all situations when the algorithm might return an invalid coloring. So, in order never to return an invalid coloring, it suffices to guard against these events.

(a) To ensure that some vertex in every initially monochromatic hyperedge h does change its color, we just need to verify that $b(u) = 1$ for some $u \in h$. This we can ensure by examining χ_0 and b just after they are generated in line 7 of Algorithm 2.

(b) To ensure that no initially non-blue hyperedge h has turned blue, we will will save the red vertices of h whenever there is potential for them all turning blue, and guard against their turning blue. Consider the hypergraph H_{Blue} with hyperedges

$$E(H_{\text{Blue}}) = \{h \cap \chi_0^{-1}(\text{Red}) : h \cap \chi_0^{-1}(\text{Red}) \neq \emptyset \text{ and } b(u) = 1 \text{ for all } u \in h \cap \chi_0^{-1}(\text{Red})\}.$$

We will then verify that no hyperedge in H_{Blue} is entirely blue in the end.

(c) To ensure that no initially non-red hyperedge h has turned red, we consider the corresponding hypergraph H_{Red} with hyperedges

$$E(H_{\text{Red}}) = \{h \cap \chi_0^{-1}(\text{Blue}) : h \cap \chi_0^{-1}(\text{Blue}) \neq \emptyset \text{ and } b(u) = 1 \text{ for all } u \in h \cap \chi_0^{-1}(\text{Blue})\},$$

and verify in the end that no hyperedge in H_{Red} is entirely red in the end.

It remains to show how H_{Blue} and H_{Red} can be stored efficiently. Each hyperedge will be stored separately by listing all its vertices. Thus, the expected sum of the sizes of the hyperedges in H_{Blue} is

$$\begin{aligned}
\mathbb{E}[\text{size}(H_{\text{Blue}})] &= \sum_{h \in E(H)} 2^{-n} \sum_{i=1}^{n} i \binom{n}{i} p^i \\
&\leq |E|\ np2^{-n} \sum_{i=0}^{n-1} \binom{n-1}{i-1} p^{i-1} \\
&\leq |E|\ np2^{-n}(1+p)^{n-1}.
\end{aligned}$$

Clearly, the same bound applies to H_{Red}. We will choose the value of p, so that with high probability $\text{size}(H_{\text{Blue}})$ and $\text{size}(H_{\text{Blue}})$ are both at most n (if either of them exceeds n, the algorithm terminates with failure), while also ensuring that the right hand sides of (4.1), (4.2) and (4.3) stay small. As before set $p = \frac{1}{2n}(\ln n - \ln\ln n)$. To the failures accounted for in (4.4), we must account for $\text{size}(H_{\text{Blue}})$ or $\text{size}(H_{\text{Blue}})$ exceeding n. Now, $\mathbb{E}[\text{size}(H_{\text{Blue}})] \leq \frac{1}{10}\sqrt{\frac{n}{\ln n}}\frac{\ln n}{2}\sqrt{\frac{n}{\ln n}} \leq \frac{n}{20}$, and by Markov's inequality $\Pr[\text{size}(H_{\text{Blue}}) \geq n] \leq \frac{1}{20}$. Similarly, $\Pr[\text{size}(H_{\text{Red}}) \geq n] \leq \frac{1}{20}$. Thus, the revised algorithm fails to deliver a valid coloring with probability at most $\frac{11}{50} + \frac{1}{20} + \frac{1}{20} \leq \frac{1}{2}$, while it uses at most $O(|V|B + nB)$ bits of space.

4.4 A local version

The above results deal with the cases when a upper bound on the number of hyperedges of H guarantees that it has a valid two-coloring. Using the Lovász Local Lemma one can show that n-uniform hypergraphs where no hyperedge intersects more than $0.17\sqrt{n/\ln n}\ 2^n$ others has a two-coloring [4]. Note that this does not require a bound on the number of hyperedges of H, but only one on the number of intersections of any one hyperedge with others. The algorithms of Moser and Tardos [10] can now be used to recover a valid two-coloring of H. As in the previous case, though the algorithm a-priori requires access to the entire hypergraph in an offline fashion, we can modify it to adapt it to the streaming setting. The details are straightforward, and we omit them here to state only the result that derives from the parallel algorithm in [10].

Proposition 4.4. *If each hyperedge of a hypergraph H intersects at most $\frac{(1-\epsilon)2^{n-1}}{e} - 1$ other hyperedges, then a two-coloring of the hypergraph can be found by a streaming algorithm that makes at most $O(\log |V|)$ passes over the input stream and uses space $O(|V|B)$.*

Compared to the Moser-Tardos algorithm, which works when the number of intersections is $O(\sqrt{n/\ln n}\ 2^n)$, this works only until intersections of up to $O(2^n)$. Furthermore, note that it requires $\log |V|$ passes over the input stream, in contrast to previous algorithms that required just one pass.

5 Coloring n-uniform hypergraphs with few vertices

The upper bound $q(n) = O(n^2 2^n)$ was shown by Erdős by considering a random graph on n^2 vertices. We ask if hypergraphs with a comparable number of hyperedges but supported on fewer vertices (perhaps $o(n)$) can be non-two-colorable. We show that that the choice of n^2 above is essentially optimal.

Theorem 5.1 (Implies Theorem 1.3(a))**.** *Let $t \geq 4$ be a parameter, and H be an n-uniform hypergraph on v vertices and q hyperedges with $v \leq \frac{n^2}{t}$ vertices. Then, a random two-coloring that partitions the vertex set into two almost equal color classes is a valid coloring with probability at least $1 - 2^{-(n-1)} \exp(-\frac{t}{8})q$.*

The following Corollary is obtained by setting $t = 8 \ln 2n$.

Corollary 5.2. *Every n-uniform hypergraph with at most $\frac{n^2}{8 \ln 2n}$ vertices and at most $n^2 2^n$ hyperedges is two-colorable.*

The Corollary shows that Erdős' construction of non-two-colorable hypergraphs having $q = n^2 2^n$ hyperedges on a vertex set of size $v = \Theta(n^2)$ is nearly optimal in terms of the number of vertices.

Proof of Theorem 5.1: Suppose $V(H) = [v]$. Consider the coloring χ obtained by partitioning $[v]$ into two roughly equal parts and coloring one part Red and the other part Blue. Every hyperedge has the same probability p of being monochromatic under this coloring. The theorem will follow immediately from the claim below using a union bound.

Claim 5.3. $p \leq 2^{-(n-1)} \exp\left(-\frac{(n-1)^2}{2v}\right)$.

Proof of claim: If $v \leq 2(n-1)$, then this probability is 0. In general, the probability that a hyperedge is monochromatic is

$$p = \frac{\binom{\lceil v/2 \rceil}{n} + \binom{\lfloor v/2 \rfloor}{n}}{\binom{v}{n}} \leq \frac{2\binom{v/2}{n}}{\binom{v}{n}} \quad \text{(since } \tbinom{a}{n} \text{ is convex for } a \geq n-1\text{)}.$$

We then have

$$\begin{aligned} p \leq \frac{2 \cdot \frac{v}{2}(\frac{v}{2}-1)(\frac{v}{2}-2)\cdots(\frac{v}{2}-n+1)}{v(v-1)(v-1)\cdots(v-n+1)} &= 2^{-(n-1)} \prod_{i=1}^{n-1} \frac{v-2i}{v-i} \\ &= 2^{-(n-1)} \prod_{i=1}^{n-1} \left(1 - \frac{i}{v-i}\right) \\ &\leq 2^{-(n-1)} \exp\left(-\sum_{i=1}^{n-1} \frac{i}{v-i}\right). \end{aligned}$$

Now,

$$\begin{aligned} -\sum_{i=1}^{n-1} \frac{i}{v-i} &= (n-1) - v\sum_{j=v-n+1}^{v-1} \tfrac{1}{j} \leq (n-1) - v\int_{v-n+1}^{v} \frac{1}{x}dx \\ &\leq (n-1) + v\ln\left(1-\frac{n-1}{v}\right) \\ &\leq (n-1) - v\left(\frac{n-1}{v} + \frac{1}{2}(\frac{n-1}{v})^2\right) \\ &= -\frac{(n-1)^2}{2v}. \end{aligned}$$

This shows that $p \leq 2^{-(n-1)}\exp\left(-\frac{(n-1)^2}{2v}\right)$ and concludes the proof of the Claim.

To recover the statement of Theorem 5.1, we union-bound over the set of hyperedges, and observe that $\frac{(n-1)^2}{v} \geq (n-1)^2 \cdot \frac{t}{n^2} \geq t/8$. □

A streaming algorithm: Theorem 5.1 (more generally, Proposition 5.5 at the end of this section for k-coloring) gives rise to a straightforward randomized streaming algorithm for two-coloring (k-coloring) assuming that the number of vertices in the hypergraph is known in advance. For example, suppose we are given an n-uniform hypergraph on $v = \frac{n^2}{t}$ vertices and at most $2^{n-2}\exp(\frac{t}{2})$ hyperedges. We just pick a random coloring χ that assigns colors Red and Blue to roughly the same number of vertices, and verify that no edge is monochromatic under χ. The algorithm uses $O(v)$ bits of space, and returns a valid coloring with probability $\frac{1}{2}$ and fails otherwise. The probability of failure can be reduced to δ by running the algorithm $\log\frac{1}{\delta}$ times in parallel.

We further examine the trade-off between the number of vertices and the number of hyperedges. Let $q(n,t)$ be the minimum number of hyperedges in a non-two-colorable hypergraph with $\frac{n^2}{t}$ vertices. Note that when the number of vertices $v \leq 2n-2$, *every* n-uniform hypergraph is trivially two-colorable by any coloring that colors $\lfloor v/2 \rfloor$ vertices Blue and the rest Red, and so we will take $\frac{n^2}{t} \geq 2n-1$. The above proof (of Theorem 1.3(a)) shows:

$$q(n,t) \geq 2^{(n-1)}\exp\left(\frac{t}{8}\right).$$

Using arguments similar to those used by Erdős [2], we will now obtain a comparable upper bound on $q(n,t)$. This makes the dependence of two-colorability on the vertex set-size explicit.

Proposition 5.4 (Restatement of Theorem 1.3(b))**.** *Suppose $N \leq \frac{n^2}{t}$ is an integer that is at least $2n-1$. Then, there is an n-uniform hypergraph with N vertices and $O(N\cdot 2^n\exp(\frac{tn}{n-2t}))$ hyperedges that is not two-colorable. Thus, $q(n,t) \leq (\frac{n^2}{t})2^n\exp(\frac{tn}{n-2t})$.*

Proof. Our non-two-colorable n-uniform hypergraph will be obtained by picking the required number of hyperedges at random from $[N]$. Fix a coloring χ. Then, a randomly drawn hyperedge is

monochromatic with probability at least

$$\begin{aligned} p = \frac{\binom{\frac{n^2}{2t}}{n}}{\binom{\frac{n^2}{t}}{n}} &= \frac{1}{2^n}\prod_{i=0}^{n-1}\frac{n^2/t-2i}{n^2/t-i} \\ &\geq 2^{-n}\left(1+\frac{t}{n-2t}\right)^{-n} \\ &\geq 2^{-n}\exp\left(\frac{-tn}{n-2t}\right). \end{aligned}$$

Let H be the hypergraph obtained by picking q hyperedges uniformly and independently from $\binom{[N]}{n}$. The probability that χ is valid for H (i.e., none of its hyperedges is monochromatic under χ) is at most

$$\left(1-2^{-n}\exp(\frac{-tn}{n-2t})\right)^q.$$

Since there are at most 2^N colorings, the probability that some coloring χ is valid for H is at most

$$2^N\left(1-2^{-n}\exp(\frac{-tn}{n-2t}\right)^q.$$

If $q \geq N \cdot 2^n \exp(\frac{tn}{n-2t}) \cdot \ln 2$ this quantity is less than 1. Our claim follows from this. □

The above results can be generalized easily to k-colorings. Let $q_k(n,t)$ be the minimum number of hyperedges in a non-k-colorable hypergraph with at most $\frac{n^2}{t}$ ($\geq kn-1$) vertices. We then have the following bounds (we omit the proofs).

Proposition 5.5. $k^{n-1}\exp\left(-\frac{(k-1)t}{2}\right) \leq q_k(n,t) \leq (\frac{n^2}{t})k^n\exp\left(\frac{(k-1)tn}{n-kt}\right)$.

6 Conclusion

The lower bound we obtain on the space requirements of one-pass streaming algorithms is optimal (up to poly(n)) factors. We present an efficient two-player two-round deterministic communication protocol for two-coloring n-uniform hypergraphs with up to $2^n/8$ hyperedges, but we do not know if there is a corresponding streaming algorithm. As mentioned in the introduction, it is easy to come up with a deterministic streaming algorithm that works in $|V|$ passes using the method of conditional expectations. Even in the two-player communication setting it would be interesting to determine if the protocol can accommodate up $\omega(2^n)$ hyperedges, perhaps even $\Omega(\sqrt{\frac{n}{\ln n}}2^n)$ hyperedges.

Our two-coloring algorithm for hypergraphs with $O(\frac{n^2}{t})$ vertices does not improve on the bound provided via the delayed recoloring algorithm when t is small, say, $o(\log n)$. We believe it should be possible to combine our argument and the delayed recoloring algorithm to show that if the number of vertices is $o(n^2)$ then we can two-color hypegraphs with strictly more than $\frac{1}{10}\sqrt{\frac{n}{\ln n}}2^n$ hyperedges.

References

[1] J.Beck. On 3-chromatic hypergraphs. *Discrete Mathematics, 24:127-137, 1978.*

[2] P. Erdős. On a Combinatorial Problem. *Nordisk Mat. Tidsskr, 11:5–10, 1963.*

[3] P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. *Colloq. Math. Soc. Jnos Bolyai, 10:609–627, 1973.*

[4] J. Radhakrishnan and A. Srinivasan. Improved bounds and algorithms for hypergraph 2-coloring. *Random Structures Algorithms, 16 (1):4–32, 2000.*

[5] S. Shannigrahi. Coloring, Embedding, Compression and Data Structure Problems on Uniform Hypergraphs *Ph.D. Thesis, School of Technology and Computer Science, Tata Institute of Fundamental Research, 2011.*

[6] N. Alon and J. H. Spencer. The Probabilistic Method. *Wiley-Interscience Series, John Wiley and Sons, Inc., New York, U.S.A., 1992.*

[7] E. Kushilevitz and N. Nisan. Communication Complexity. *Book, Cambridge University Press, 1996.*

[8] N. Alon, Y. Matias and M. Szegedy. The space complexity of approximating the frequency moments. *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing (STOC '96) Pages 20-29, 1996.*

[9] D. Achlioptas, J.H. Kim, M. Krivelevich, and P. Tetali. Two-coloring random hypergraphs. *Random Structures & Algorithms* 20(2): 249-259, 2002.

[10] R. A. Moser and G. Tardos. A constructive proof of the general Lovász local lemma. *Journal of the ACM* 57 (2): 1-15, 2010.

[11] M. Mitzenmacher and E. Upfal. Probability and Computing: Randomized Algorithms and Probabilistic Analysis. Cambridge University Press, 2005.

[12] R. M. Karp, C. H. Papadimitriou and S. Shenker. A simple algorithm for finding frequent elements in streams and bags. *ACM Transactions on Database Systems* 28: 51-55, 2003.

[13] D. D. Cherkashin and J. Kozik. A note on random greedy coloring of uniform hypergraphs. *Random Structures & Algorithms* 47(3): 407-413, 2015.

[14] M. Karchmer and A. Wigderson. Monotone Circuits for Connectivity Require Super-Logarithmic Depth. *SIAM J. Discrete Math.* 3(2): 255-265 , 1990.

[15] M. Ajtai. Σ_1^1-formulae on nite structures.*Ann. Pure Appl. Logic.* 24(1):148, 1983.

[16] Emanuele Viola. On Approximate Majority and Probabilistic Time. *Computational Complexity*, v.18 n.3, p.337-375, 2009.

[17] S. Muthukrishnan. *Data Streams: Algorithms and Applications*, Foundations and Trends in Theoretical Computer Science: Vol. 1: No. 2, pp 117-236, 2005.

[18] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, Jian Zhang. *On Graph Problems in a Semi-streaming Model*, Proc. 31st International Colloquium on Automata, Languages and Programming, (ICALP), v. 3182: 531-543, 2004.

[19] J. Feigenbaum, S. Kannan, and J. Ziang. *Computing diameter in the streaming and sliding window models*, Algorithmica, vol. 41, no. 1, pp. 2541, 2004

[20] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. *Graph distances in the streaming model: The value of space*. Proc. ACM-SIAM SODA, pp. 745754, 2005.